



Using Text Injection to Improve Recognition of Personal Identifiers in Speech

Yochai Blau, Rohan Agrawal, Lior Madmony, Gary Wang, Andrew Rosenberg, Zhehuai Chen, Zorik Gekhman, Genady Beryozkin, Parisa Haghani, Bhuvana Ramabhadran

Google

{yochaib, rohanag, liomad, wgary, rosenberg, zhehuai, zorik, genady, parisah, bhuv}@google.com

Abstract

Accurate recognition of specific categories, such as persons' names, dates or other identifiers is critical in many Automatic Speech Recognition (ASR) applications. As these categories represent personal information, ethical use of this data including collection, transcription, training and evaluation demands special care. One way of ensuring the security and privacy of individuals is to redact or eliminate Personally Identifiable Information (PII) from collection altogether. However, this results in ASR models that tend to have lower recognition accuracy of these categories. We use text-injection to improve the recognition of PII categories by including fake textual substitutes of PII categories in the training data using a text injection method. We demonstrate substantial improvement to Recall of Names and Dates in medical notes while improving overall WER. For alphanumeric digit sequences we show improvements to Character Error Rate and Sentence Accuracy.

Index Terms: conformers, E2E models, Medical ASR, Text-Injection Training, protecting PHI, de-identification, PII

1. Introduction

While a lot of speech is publicly broadcast or shared online to a vast audience, speech recognition applications frequently interact with more private communications such as dictation, call-center conversations or conversing with a digital assistant. Automatic speech recognition (ASR) systems are trained using transcribed speech, and operate best when their training data matches the context in which they are used. This raises a question of how to train ASR to work well on private communications, while being sensitive to the complexities of collecting and transcribing training material from private contexts.

We engage with this question in the context of healthcare. Accurate transcription of medical speech is key to a growing set of applications, including dictation of clinical notes and voice assistance. However, ethical collection and transcription of training data for medical ASR requires care. Medical data is exceptionally sensitive and respecting patients' privacy in any collection effort is essential. De-identification, a process which includes the redaction of PII tokens, is a common practice to protect users' private data. In the process, the audio segment is replaced by silence and the text is either removed or replaced by a special markup tag. De-identification is used extensively for healthcare data and is often required by regulations such as the US HIPAA Act [1].

In contexts such as call-center applications, speech utterances are shorter and may be in response to prompts like "Please confirm your date of birth" or "Please say your user id". In this case, de-identification results in completely eliminating the utterance. In both domains, using such de-identified datasets for

training ASR models makes recognition of these classes of PII terms operate at a higher error rate than the surrounding speech.

Prior work has focused around replacing PII in speech with arbitrary synthesized speech from the same category [2] (for example, named entities, number sequences) or ensuring that these PII tokens/sequences are adequately covered by a language model, as done in [3]. Text injection has recently emerged as an effective way to leverage text-only data for ASR training in a single model [4, 5, 6, 7] without an LM. Techniques such as the one described in [4] have shown promise for zero-transcribed speech adaptation. In [8], Wang & Kastner et al. demonstrated that text-injection can effectively address domain transfer on diverse SpeechStew corpora [9] with little to no in-domain speech. Motivated by this, we propose to generate a text dataset that contains fake terms instead of the redacted PII and use text-injection during model training. This eliminates the need for any sensitive transcribed training data and allows effective use of de-identified medical datasets.

This paper addresses three specific data challenges when training a model for private contexts with the use of text-injection in ASR training.

- Domain adaptation: Transcribed speech from private contexts is much harder to come by than public contexts. We show that text-injection along with a small amount of transcribed speech can effectively perform domain adaptation to Medical ASR resulting in a 11.5 reduction to WER and improvements of 1% to Names and 11% to Dates (Section 5.1).
- Redaction: Available in-domain data has redacted PII from the transcript with corresponding speech replaced by silence. To mitigate this, we train on generated, spoken text similar to the PII, which allows for substantial improvement to sensitive term recognition without using any individual's protected data; compared to training on redacted speech data, Names and Dates recall is improved by 8% and 13% (Section 5.1).
- Elimination: De-identification may eliminate utterances that contain sensitive terms. For example, short utterances containing alphanumeric and digit sequences are eliminated. We show that text-injection can substantially improve recognition of these classes of potentially sensitive information with a full-sequence Sentence Accuracy improvement of 3.2% (and CER improvement of 1.4%) (Section 5.2).

2. Related Work

The complexity of recognizing personal identifiers in medical speech has recently been studied in [2]. This work shows that PII recognition accuracy declines when training an end-to-end model on de-identified data. The authors propose to generate artificial audio/text pairs with synthetic identifiers, and show that training on such data restores most of the performance degrada-

tion. Artificial audio/text pairs are generated in a two-step process: first, PII in the text is replaced with random data. Then, corresponding audio is generated or spliced with matching fragments. Generating or manipulating the audio is more demanding than replacing identifiers in the text and also requires accurate word-level timestamps. In our work we take this concept one step further, by improving PII recognition with text replacements alone, and *without* any audio generation or manipulation. Other work has studied medical domain ASR [3, 10, 11, 12], yet these have not reported findings on PII recognition accuracy and the effect of training on de-identified data.

Regarding text-injection in ASR training, this work builds on the architecture described in [4]. This is described in detail in Section 3. However, there is a set of related approaches and architectures that similarly use separate speech and text encoders that feed into a shared encoder [5, 6, 7]. Additionally, a similar thread of work trains on speech and text data operates by first converting the text data to speech via TTS [13, 14, 15, 16, 17]. These approaches, in effect, use speech as an intermediate representation to perform text injection.

3. Text-Injection and ASR Details

The text-injection model used in this work includes a speech encoder, a text encoder with a learned duration model, a shared encoder, decoder and an alignment decoder following [4]. The speech encoder consists of 4 causal conformer layers. The shared encoder consists of 3 causal conformer layers and 10 non-causal conformer layers, each with model dimension of 512. The text encoder contains 2 conformer layers and 4 lightweight convolutional upsampling layers [18]. HAT decoders [19] with v^2 embeddings [20] are used in both decoder and an alignment decoder with the distinction of the former produces word-piece models as text outputs and the latter uses phonemes as model units to get speech-text alignments. The overall model contains 165M parameters, with an additional 58M in the text encoder which is only used during training.

Text-injection training involves speech-text and text-only training paths in a curriculum fashion. Initially, speech-text training is used to minimize a consistency loss, a HAT decoder loss and an alignment decoder loss, for the duration model. Consistency loss ensures that we learn corresponding mappings from text to speech embedding. The duration model is used to up-sample text embedding before being fed to the shared encoder. After beginning training with speech-text loss, we enable text-only loss to be able to perform text-injection training. The text-only training step involves minimizing an aligned masked-language model loss. Further details of text injection architecture and training can be found in [4].

In all experiments, we train a text-injection model on supervised only data for 10k steps first so that the speech and text encoders generate consistent features. We continue training with target domain text or targeted terms as described in the following sections. While training with text alone is feasible after the initial supervised training, we find that including supervised data during training helps to stabilize the model behavior.

4. Data Description

We are aware of the sensitive nature of speech recognition research particularly on personal identifiers. Therefore, we ensure that this work abides by the Google AI Principles [21].

4.1. Medical domain datasets

The *Medical Audio* dataset consists of dictations of clinical notes by healthcare professionals from a variety of medical specialties. During the transcription process any potentially identifying information was redacted. The corresponding audio segment was replaced by silence and a special markup was placed in the transcript to indicate the type of the data that was removed. Examples of the markup tags include `PATIENT_NAME`, `MEDICAL_PROFESSIONAL_NAME`, `AGE` and `DATE`. The resulting dataset was certified to comply with de-identification requirements of the US HIPAA Privacy Rule [1].

The *Medical Text* dataset consists of the transcriptions from the *Medical Audio* dataset (without the audio recordings). As mentioned above, these transcriptions include markup tags for PII's such as names, addresses, dates, etc. that were redacted. In this dataset, we replace four types of markup tags (names, identifying numbers, dates and ages) with fake random data. For example, any name tag is replaced by a random name from a real-world distribution, each redacted digit is replaced with a random digit between 0-9, and so on. Note that replacing redacted PII information in text is quite easy and straightforward. However, splicing synthetic speech within speech recordings where the PII was silenced in a natural manner is challenging and error-prone, especially in the case of fast-paced speech recorded in a noisy environment as it is in the *Medical Audio* dataset.

The *Synthetic Notes* dataset is a collection of fake dictations, recorded by clinicians in a noise-free environment and consists of 14.5 hours of audio. The dataset consists of synthetic hospital visit summaries which include fake identifiers such as names, IDs, dates, etc. The dataset creation process started with creation of fake hospital visit summaries by trained clinicians. Later, a separate group of clinicians simulated dictations of hospital notes based on these visit summaries, which were then independently transcribed.

The *Synthetic Names* dataset is a 14-hours long dataset that consists of short phrases (average length 7.6s) that resemble phrases in the dictation of clinical notes and include references to patient or clinician names. First, medical specialists created 550 textual templates with placeholders for names, ensuring diversity of name formats. Each template was recorded in a quiet environment by 17 different speakers, with placeholders substituted by names, each time randomly drawn from a real-world distribution. All name occurrences are tagged, which allows to compute a recall metric for names alone (along with more global metrics such as WER).

Table 1: *Dataset descriptions.*

Name	Size (hours)	Type	Description
Training data			
Captions	377K	Speech	Speech from videos and corresponding captions.
Medical Audio	5K	Speech	Clinicians dictating notes. All private information is redacted.
Medical Text	-	Text	The transcriptions from the <i>Medical Audio</i> dataset, where redacted identifiers are replaced with random data.
Testing data			
Synthetic Notes	14.5	Speech	Synthetic hospital notes dictated by clinicians, including (fake) identifiers.
Synthetic Names	14	Speech	Synthetic short phrases from clinical notes, including (fake) names (patient/clinician/family member/etc.).

4.2. Alphanumeric Sequence Identifier Datasets

For personal identifier text-injection train data, we use 100M sequences each of alphanumeric sequences and text sequences. The length of these sequences is sampled from a skewed normal distribution $X(\mu = 10, \sigma = 5) + 3$. For digit sequences, characters from 0-9 are sampled uniformly. For alphanumeric sequences, characters from the alphanumeric set 0-9,a-z are sampled uniformly. 10% of sequences are chosen at random, and character repeats of length 2, 3 or 4 are added. This is done because character or digit repetitions are a challenging use case for an ASR system.

The *spoken digit sequence* test set consists of 2146 utterances (average length 4.3s) with randomly generated digit sequences, spoken by voice actors.

The *spoken alphanumeric sequence* test set consists of 1,992 utterances (average length 5.2s) with randomly generated alphanumeric sequences, spoken by voice actors. The voice actors in both sequence test sets have acted out hesitations and pauses while speaking these sequences.

The following TTS datasets contain utterances synthesized by a commercial American English TTS system sampled from 6 voices. In this work, we use the TTS datasets as test data, and the spoken sequence data as either test or training material (Section 5.2).

The *TTS digit sequence* test set consists of 1991 utterances (average length 5.4s) with randomly generated digit sequences.

The *TTS alphanumeric sequence* test set consists of 2000 utterances (average length 5.7s) with randomly generated alphanumeric sequences.

The *TTS digit sequence repetition* consists of 3188 utterances (average length 5.4s) with randomly generated digit sequences which have at minimum, 1 consecutive digit repetition. 1001 utterances have 2, 3 and 4 consecutive repetitions and 185 utterances have more than 4 repetitions.

The *TTS alphanumeric sequence repetition* consists of 2179 utterances (average length 5.7s) with randomly generated alphanumeric sequences which have at minimum, 1 consecutive character repetition. 1001 utterances have 2, 3 consecutive repetitions and 177 utterances have more than 4 repetitions.

4.3. General datasets

The base model (B1) is trained with close to 300k hours of multidomain utterances including YouTube, Telephony and Dictation for US English as described in [22]. Multi-condition training [23], random 8kHz down-sampling [24] are applied on the training data. Importantly, any utterance that includes alphanumeric and digit sequences are eliminated from transcribed training data, as they may represent personal identifiers. This data is only included in initial checkpoint training, described as A1/B1 in Section 5, and is not further used during finetuning.

Along with the medical domain data described above, we also include 377k hours of YouTube *Captions* data as described in [25] to maintain strong general purpose ASR performance. While this material may incidentally include medical content, no particular targeting, selection or filtering for in-domain material was performed.

5. Recognizing Redacted and Eliminated Terms

5.1. Identifiers in medical speech

All models used the architecture described in Section 3, and use a base checkpoint, B1, trained on multidomain data (Section 4.3). The A1 model was trained on the (non-medical) *Captions* dataset, in which medical terms are relatively uncommon. This shows clear domain-transfer effects; the WER on the medical-term rich *Synthetic Notes* dataset is high, with a WER of 14.6% (Table 2).

Adapting this model to the medical domain is done by training on a mixture of *Captions* and *Medical Audio* datasets with a 90%/10% ratio, resulting in model A2. The *Synthetic Notes* WER drops significantly, and we observe higher recall for medical entities such as conditions and medications. Yet, for other entities such as names and dates the recall metrics degrade. This is a direct consequence of names, dates and other identifiers being redacted from the *Medical Audio* dataset.

To improve error rates on PII, model T1 is trained on the same mixture of speech data as for A2, with the *Medical Text* dataset injected during training. This textual data contains identifiers such as names and dates, leading to a 8%/13% boost in names/dates recall, respectively. In Table 3 we show two cases where model T1, which has been exposed to medical text with (fake) identifiers, performs better in transcribing examples from the *Synthetic Names* dataset. We also observe that the introduction of medical text further improves the overall WER from 3.3 to 3.1, demonstrating the value of additional in-domain text to address domain transfer as well as targeted terms.

Table 2: Word error rate for *Synthetic Notes* and medical-term/name/date recall for *Synthetic Names*.

	Datasets	Synthetic Notes WER	Medical term recall	Names recall	Dates recall
B1	Multidomain (MD)	14.6	82.6%	60%	66%
A1	MD, Captions	14.7	82.1%	59%	72%
A2	MD, Captions Medical Audio	3.3	97.4%	53%	64%
T1	MD, Captions Medical Audio Medical Text (injected)	3.1	97.5%	61%	77%

Table 3: Examples of name mistranscriptions from *Synthetic Names corpus*. Note: All examples are fake descriptions of hypothetical scenarios.

Model	Transcription
Truth	Scarlett Kathleen Ibarra was admitted to the floors
A2	Scarlet Caffeine Ebara was admitted to the floors
T1	Scarlett Kathleen Ibarra was admitted to the floors
Truth	Oliver Barry Matthews is adamant that the patient had no past surgeries
A2	All of our bearing math uses adamant that the patient had no past surgeries
T1	Oliver Barry Matthews is adamant that the patient had no past surgeries

The training data used for models A1, A2 and T1 above, includes a large general-domain multidomain and *Captions* dataset. This dataset includes identifying entities such as names and dates to some extent. To isolate the effect of PII redaction, we train a cold-start model A3 only on the *Medical Audio* dataset (where all identifiers are redacted). The results of this experiment are in Table 4. All error metrics degrade, but particularly noticeable is the names recall which drops to 1%, as

the model has not seen *any* names during training. Yet performance can be improved by injecting text which includes fake identifiers, even without a single audio example of such entities, as seen for model T2. While the metrics indicate that text injection alone cannot match the performance of T1, the recall of Names and Dates are improved by a factor of 1800% and 207%, respectively. Moreover, the introduction of medical text improves the WER performance of the model from 6.1 to 4.1.

Table 4: Ablation study: training only on medical data.

	Datasets	Synthetic Notes WER	Medical term recall	Names recall	Dates recall
A3	Medical Audio	6.1	96.2%	1%	28%
T2	Medical Audio Medical Text (injected)	4.1	97.2%	18%	58%

5.2. Recognizing Alphanumeric Sequence Identifiers

The architecture of the base ASR model B1 is the same RNN-T model with a cascade encoder and HAT decoder with v^2 embeddings in Section 5.1. The baseline text injection model, M1 was trained with the same paired text data as B1, and additional base text injection data (not related to target domain) consists of 100B anonymized sentences across different domains. Note that this text injection model is trained in arbitrary text to improve core ASR performance rather than targeting any particular domain or type of lexical content.

We aim to improve recognition of alphanumeric and digit sequence identifiers. Thus, we perform a round of text injection training based on the initial model described above, specifically, the personal identifier text data (cf. Section 4.2). Results of these experiments measured by character error rate (CER) and Sentence Accuracy (SACC) in Tables 5 and 6. Character error rate is more fine grained, by being able to measure individual errors. However, when recognizing an identifier it is crucial to recognize the full sequence; this is measured by SACC.

Table 5: CER on alphanumeric and digit test sets.

Dataset	B1: Base model	M1: B1 + base text injection	M2: M1 + in-domain text inj.	M3: M2+ Spoken datasets
Spoken digit	3.1	2.4	2	-
TTS digit rep.	15.2	15.7	14.5	17.0
TTS digit	1.7	1	1.1	0.8
Spoken alphanumeric	8.6	7.3	6.8	-
TTS alphanumeric rep.	12.1	11.9	9.3	8.4
TTS alphanumeric	9.2	8.5	8	6.2
average CER	8.3	7.8	6.9	-

Table 6: SACC on alphanumeric and digit test sets.

Dataset	B1: Base model	M1: B1 + base text injection	M2: M1 + in-domain text inj.	M3: M2+ Spoken datasets
Spoken digit	93.8	94.7	95.2	-
TTS digit rep.	67.4	69.1	70.4	64.3
TTS digit	96.8	96.6	97.1	96.8
Spoken alphanumeric	73.6	75.5	76.5	-
TTS alphanumeric rep.	54.2	53.5	63.9	63.6
TTS alphanumeric	66.0	67.4	68.1	72.4
average SACC	75.3	76.1	78.5	-

We use CER as a metric instead of word error rate for evaluation of personal identifier recognition because CER is a better fit for these sequences; an utterance may contain a single “word” like *A1B2C3*. Normalization helps to ignore formatting

issues with respect to such sequences. As can be seen from Table 5, even the addition of base text injection data which is not related to target domain (M1) improves CER from 8.3 to 7.8, a reduction of 6% from B1. Additional training on domain specific text injection data (M2), further reduces normalized CER to 6.9, a reduction of 14.1% from M1. Correspondingly when considering SACC, after adding base text injection data (M1), we see an increase from 75.3 to 76.1, a modest increase of 1.1% from B1. SACC further increases to 78.5 after training on domain specific text injection data (M2), an increase of 3.1% from M1. In M3, we add Spoken sequence data (4138 utterances) to our training corpus. The motivation for this experiment is to show the effect of adding a small amount of paired speech-text data in addition to text-injection training can further help. This further reduces CER on alphanumeric test sets, but gives mixed results in digit test sets.

Catastrophic forgetting is a concern when fine-tuning a pre-trained model on material from a narrow domain, as we are doing here with text-injection of sequences. To assess the impact of this we assess performance on a Short Utterance test set representative of spoken search queries, and TTS test sets focusing on rare proper names from 5 different domains (Table 7). While we see a minor regression between M1 and M2, the performance is still substantially better than B1, the base model with no text-injection. This suggests that this regression can be mitigated with tuning of a ratio and schedule of the text-injection training.

Table 7: WER on broad domain test sets.

Dataset	B1: Base model	M1: B1 + base text injection	M2: M1 + in-domain text inj.
Short Utterances	6.7	5.6	5.7
TTS rare terms 1	14.5	11.3	11.9
TTS rare terms 2	9.8	8.6	8.9
TTS rare terms 3	37.9	33.5	34.4
TTS rare terms 4	21.9	17.9	18.8
TTS rare terms 5	24.9	20.7	21.4

6. Conclusions

Due to the sensitive nature of medical speech and personal identifiers, substantial care must be taken in the collection, transcription and use of this data, often requiring the removal of personal identifiers from the data. Models trained on such data have higher error rates when recognizing such identifiers. In this paper, we demonstrate that text-injection is an effective approach to improving ASR performance of models trained on de-identified speech, including improved recognition of personal identifiers. Text-injection can leverage text where speech is unavailable, and synthetic text is simple to construct, frequently indistinguishable from real transcripts, and not associated with any person’s private communication. In Medical ASR where names and dates are redacted from speech and transcripts, we find core WER improves from 3.3 to 3.1, and recall of names and dates are improved by 8% and 13% respectively. For short utterances containing alphanumeric and digit sequences which can be used as personal identifiers in a broad range of contexts, we improve SACC by an average of 3.2% while CER is reduced from 8.3% to 6.9%

7. References

- [1] “The Health Insurance Portability and Accountability Act of 1996 (HIPAA),” <https://www.hhs.gov/hipaa/index.html>, 1996.
- [2] M. Flechl, S.-C. Yin, J. Park, and P. Skala, “End-to-end speech

- recognition modeling from de-identified data,” in *Proc. Interspeech*, 2022, pp. 1382–1386.
- [3] E. Edwards, W. Salloum, G. P. Finley, J. Fone, G. Cardiff, M. Miller, and D. Suendermann-Oeft, “Medical speech recognition: Reaching parity with humans,” in *Speech and Computer*, 2017, pp. 512–524.
 - [4] Z. Chen, Y. Zhang, A. Rosenberg, B. Ramabhadran, P. J. Moreno, A. Bapna, and H. Zen, “MAESTRO: Matched Speech Text Representations through Modality Matching,” in *Proc. Interspeech*, 2022, pp. 4093–4097.
 - [5] A. Bapna, C. Cherry, Y. Zhang, Y. Jia, M. Johnson, Y. Cheng, S. Khanuja, J. Riesa, and A. Conneau, “mslam: Massively multilingual joint pre-training for speech and text,” *arXiv preprint arXiv:2202.01374*, 2022.
 - [6] T. N. Sainath, R. Prabhavalkar, A. Bapna, Y. Zhang, Z. Huo, Z. Chen, B. Li, W. Wang, and T. Strohmman, “Joist: A joint speech and text streaming model for asr,” in *2022 IEEE Spoken Language Technology Workshop (SLT)*, pp. 52–59.
 - [7] S. Thomas, B. Kingsbury, G. Saon, and H.-K. J. Kuo, “Integrating text inputs for training and adapting rnn transducer asr models,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 8127–8131.
 - [8] G. Wang, K. Kastner, A. Bapna, Z. Chen, A. Rosenberg, B. Ramabhadran, and Y. Zhang, “Understanding shared speech-text representations,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
 - [9] W. Chan, D. Park, C. Lee, Y. Zhang, Q. Le, and M. Norouzi, “Speechstew: Simply mix all available speech recognition data to train one large neural network,” *arXiv preprint arXiv:2104.02133*, 2021.
 - [10] C.-C. Chiu, A. Tripathi, K. Chou, C. Co, N. Jaitly, D. Jaunzeikare, A. Kannan, P. Nguyen, H. Sak, A. Sankar, J. Tansuwan, N. Wan, Y. Wu, and X. Zhang, “Speech Recognition for Medical Conversations,” in *Proc. Interspeech*, 2018, pp. 2972–2976.
 - [11] A. Mani, S. Palaskar, N. V. Meripo, S. Konam, and F. Metzger, “Asr error correction and domain adaptation using machine translation,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6344–6348.
 - [12] L. Zhou, S. V. Blackley, L. Kowalski, R. Doan, W. W. Acker, A. B. Landman, E. Kontrient, D. Mack, M. Meteer, D. W. Bates, and F. R. Goss, “Analysis of errors in dictated clinical documents assisted by speech recognition software and professional transcriptionists,” *JAMA network open*, vol. 1, no. 3, pp. e180530–e180530, 2018.
 - [13] A. Rosenberg, Y. Zhang, B. Ramabhadran, Y. Jia, P. Moreno, Y. Wu, and Z. Wu, “Speech recognition with augmented synthesized speech,” in *IEEE automatic speech recognition and understanding workshop (ASRU)*, 2019, pp. 996–1002.
 - [14] Z. Chen, Y. Zhang, A. Rosenberg, B. Ramabhadran, P. Moreno, and G. Wang, “Tts4pretrain 2.0: Advancing the use of text and speech in asr pretraining with consistency and contrastive losses,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.
 - [15] A. Fazel, W. Yang, Y. Liu, R. Barra-Chicote, Y. Meng, R. Maas, and J. Droppo, “SynthASR: Unlocking Synthetic Data for Speech Recognition,” in *Proc. Interspeech*, 2021, pp. 896–900.
 - [16] X. Zheng, Y. Liu, D. Gunceler, and D. Willett, “Using synthetic audio to improve the recognition of out-of-vocabulary words in end-to-end asr systems,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 5674–5678.
 - [17] G. Wang, A. Rosenberg, Z. Chen, Y. Zhang, B. Ramabhadran, Y. Wu, and P. Moreno, “Improving speech recognition using consistent predictions on synthesized speech,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7029–7033.
 - [18] I. Elias, H. Zen, J. Shen, Y. Zhang, Y. Jia, R. J. Weiss, and Y. Wu, “Parallel tacotron: Non-autoregressive and controllable tts,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 5709–5713.
 - [19] E. Variani, D. Rybach, C. Allauzen, and M. Riley, “Hybrid autoregressive transducer (hat),” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
 - [20] M. Ghodsi, X. Liu, J. Apfel, R. Cabrera, and E. Weinstein, “Rnn-transducer with stateless prediction network,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7049–7053.
 - [21] “Our principles - Google AI,” <https://ai.google/principles/>, 2023, [Online; accessed 01-March-2023].
 - [22] A. Narayanan, R. Prabhavalkar, C.-C. Chiu, D. Rybach, T. N. Sainath, and T. Strohmman, “Recognizing long-form speech using streaming end-to-end models,” in *IEEE automatic speech recognition and understanding workshop (ASRU)*, 2019, pp. 920–927.
 - [23] C. Kim, A. Misra, K. Chin, T. Hughes, A. Narayanan, T. N. Sainath, and M. Bacchiani, “Generation of Large-Scale Simulated Utterances in Virtual Rooms to Train Deep-Neural Networks for Far-Field Speech Recognition in Google Home,” in *Proc. Interspeech*, 2017, pp. 379–383.
 - [24] J. Li, D. Yu, J.-T. Huang, and Y. Gong, “Improving wide-band speech recognition using mixed-bandwidth training data in cd-dnn-hmm,” in *IEEE Spoken Language Technology Workshop (SLT)*, 2012, pp. 131–136.
 - [25] H. Liao, E. McDermott, and A. Senior, “Large scale deep neural network acoustic modeling with semi-supervised training data for youtube video transcription,” in *IEEE Workshop on Automatic Speech Recognition and Understanding*, 2013, pp. 368–373.